

Applicant : Gadir et al.
Serial No. : 09/911,902
Filed : July 23, 2001
Page : 2 of 15

Attorney's Docket No.: 07575-053001

Amendments to the Claims:

This listing of claims replaces all prior versions and listings of claims in the application:

Listing of Claims:

1. (Currently Amended) A file server system, comprising:
two or more nodes, each node configured to run two or more virtual servers, each virtual server having as exclusive resources a virtual interface to clients and one or more file systems, wherein an exclusive resource of a virtual server can be seen by the virtual server and not by another virtual server.
2. (Original) The system of claim 1, wherein the virtual interface comprises a virtual IP address.
3. (Original) The system of claim 1, wherein the virtual interface comprises two or more virtual IP addresses.
4. (Original) The system of claim 1, wherein clients access the file systems using NFS or CIFS protocols.
5. (Original) The system of claim 1, further comprising failover computer program instructions operable to be executed to cause the system to:
detect a failure of a first node; and
migrate each virtual server on the first node to a different node in the system.

Applicant : Gadir et al.
Serial No. : 09/911,902
Filed : July 23, 2001
Page : 3 of 15

Attorney's Docket No.: 07575-053001

6. (Original) The system of claim 5, wherein each virtual server has an associated failover priority, and the failover instructions further comprise instructions to:
migrate virtual servers in order of their respective priorities.

7. (Original) The system of claim 5, wherein the failover instructions further comprise instructions to:

recognize a virtual server that is identified as not to be migrated in the event of node failure and prevent migration of a so-identified virtual servers when it is on a node that fails.

8. (Original) The system of claim 1, further comprising rerouting computer program instructions operable to be executed to cause the system to:

detect a failure in a first subnet connected to a first node, the first node having a network connection to a first client;

identify a second node having a network connection to the first client and a connection over a second, different subnet to the first node;

use the second node as a router in response to the detected failure to route data between the first client and the first node.

9. (Original) The system of claim 8, wherein before failure in the first subnet, the connection between the first client and the first node is through a first virtual IP address assigned to a port on the first node, the rerouting instructions further comprising instructions to:

migrate the first virtual IP address to a port on the second node connected to the second subnet.

Applicant : Gadir et al.
Serial No. : 09/911,902
Filed : July 23, 2001
Page : 4 of 15

Attorney's Docket No.: 07575-053001

10. (Original) The system of claim 1, further comprising failover computer program instructions operable to be executed to cause the system to:

- detect a failure of a physical port on a first node;
- determine whether any other physical port on the first node is good;
- migrate all virtual IP addresses associated with the failed physical port to a good physical port on the first node if there is such a good port; and
- migrate all virtual IP addresses associated with the failed physical port along with all virtual servers attached to such virtual IP addresses to a different, second node if there is no such good port on the first node.

11. (Original) The system of claim 10, wherein the failed physical port is on a first subnet and the good physical port is on a different, second subnet.

12. (Original) The system of claim 1, wherein the system comprises load-balancing computer program instructions operable to be executed to cause the system to:

- calculate a balanced distribution of the virtual server loads across the nodes of the system, excluding any failed nodes; and
- perform load balancing by migrating one or more virtual servers from heavily loaded nodes to less heavily loaded nodes.

13. (Original) The system of claim 1, further comprising computer program instructions operable to be executed on a first node to:

- determine a load on each physical port on the first node; and
- redistribute the virtual interfaces on the first node among the physical ports of the first node for load balancing over the physical ports.

Applicant : Gadir et al.
Serial No. : 09/911,902
Filed : July 23, 2001
Page : 5 of 15

Attorney's Docket No.: 07575-053001

14. (Original) The system of claim 1, further comprising computer program instructions operable to be executed to cause the system to:

detect an inability on a first node to access of shared storage unit; and

in response to detection of the inability to access the shared storage unit, migrate all virtual servers containing file systems on the shared storage unit to an alternative node that can access the storage unit if such an alternative node exists in the system.

15. (Original) The system of claim 12, wherein the load-balancing instructions are further operable to determine a load on each virtual server.

16. (Original) The system of claim 12, wherein the load-balancing instructions are further operable to determine a load on each physical server.

17. (Original) The system of claim 12, wherein the nodes include a master node and the load-balancing instructions are operable to be executed on the master node.

18. (Original) The system of claim 12, wherein the load-balancing instructions are operable to migrate a first virtual server and a second virtual server from a first node, the first virtual server being migrated to a second node of the system and the second virtual server being migrated to a different, third node of the system.

19. (Original) The system of claim 12, wherein the load-balancing instructions are operable to balance system load as part of a failover process.

20. (Original) The system of claim 12, wherein the load-balancing instructions are operable to balance system load independent of any failover occurring.

Applicant : Gadir et al.
Serial No. : 09/911,902
Filed : July 23, 2001
Page : 6 of 15

Attorney's Docket No.: 07575-053001

21. (Original) The system of claim 1, further comprising computer program instructions operable to be executed to cause the system to:

detect without user intervention a file system degradation of a first file system; and
block access to the first file system in response to the detection of the degradation, repair the first file system, and then permit access to the first file system, all without user intervention.

22. (Currently Amended) A file server system, comprising:

a node configured with a virtual server having two or more simultaneously active virtual IP addresses, the virtual server having as exclusive resources one or more file systems, wherein an exclusive resource of a virtual server can be seen by the virtual server and not by another virtual server.

23. (Original) The system of claim 22, wherein the node is configured with a second virtual server having two or more other simultaneously active virtual IP addresses.

24. (Currently Amended) A file server system, comprising:

two or more nodes, each node being configured to run a virtual server having as exclusive resources a virtual IP address and one or more file systems, wherein an exclusive resource of the virtual server can be seen by the virtual server and not by another virtual server, and each node being configured with two or more physical ports;

wherein a first node is further configured to:

detect a failure of a physical port on the first node;
determine whether any other physical port on the first node is good;
migrate all virtual IP addresses associated with the failed physical port to a good physical port on the first node if there is such a good port; and
migrate all virtual IP addresses associated with the failed physical port along with all virtual servers attached to such virtual IP addresses to a different, second node if there is no such good port on the first node.

Applicant : Gadir et al.
Serial No. : 09/911,902
Filed : July 23, 2001
Page : 7 of 15

Attorney's Docket No.: 07575-053001

25. (Currently Amended) A computer program product, tangibly stored on a computer-readable medium or propagated signal, for execution in multiple nodes of a file server system cluster, comprising instructions operable to cause a programmable processor to:

- detect a failure of a first node of the cluster; and
- migrate each of multiple virtual servers on the first node to a different node in the cluster each virtual server having as exclusive resources one or more file systems, wherein an exclusive resource of a virtual server can be seen by the virtual server and not by another virtual server.

26. (Original) The product of claim 25, further comprising instructions to:
migrate virtual servers in order of their respective priorities.

27. (Previously Presented) The product of claim 25, further comprising instructions to:
recognize a virtual server that is identified as not to be migrated in the event of node failure and prevent migration of a so-identified virtual server when it is on a node that fails.

28. (Original) The product of claim 25, further comprising instructions to:
detect a failure in a first subnet connected to a first node, the first node having a network connection to a first client;
identify a second node having a network connection to the first client and a connection over a second, different subnet to the first node;
use the second node as a router in response to the detected failure to route data between the first client and the first node.

Applicant : Gadir et al.
Serial No. : 09/911,902
Filed : July 23, 2001
Page : 8 of 15

Attorney's Docket No.: 07575-053001

29. (Original) The product of claim 25, further comprising instructions to:
detect a failure of a physical port on a first node of the cluster;
determine whether any other physical port on the first node is good;
migrate all virtual IP addresses associated with the failed physical port to a good physical port on the first node if there is such a good port; and
migrate all virtual IP addresses associated with the failed physical port along with all virtual servers attached to such virtual IP addresses to a different, second node of the cluster if there is no such good port on the first node.

30. (Original) The product of claim 29, wherein before failure in the first subnet, the connection between the first client and the first node is through a first virtual IP address assigned to a port on the first node, the rerouting instructions further comprising instructions to:
migrate the first virtual IP address to a port on the second node connected to the second subnet.

31. (Original) The product of claim 25, further comprising load-balancing instructions to:
determine a load produced by each virtual server;
calculate a balanced distribution of the virtual server loads across the nodes of the server, excluding any failed nodes; and
perform load balancing by migrating one or more virtual servers from heavily loaded nodes to less heavily loaded nodes.

32. (Original) The system of claim 31, wherein the nodes include a master node and the load-balancing instructions are operable to be executed on the master node.

Applicant : Gadir et al.
Serial No. : 09/911,902
Filed : July 23, 2001
Page : 9 of 15

Attorney's Docket No.: 07575-053001

33. (Original) The system of claim 31, wherein the load-balancing instructions are operable to migrate a first virtual server and a second virtual server from a first node, the first virtual server being migrated to a second node of the system and the second virtual server being migrated to a different, third node of the system.

34. (Currently Amended) A computer program product, tangibly stored on a computer-readable medium or propagated signal, for execution in a node of a file server system cluster ~~in which virtual servers have virtual IP addresses associated with physical ports~~, the product comprising instructions operable to cause a programmable processor to:

detect a failure of a physical port on a first node of the cluster, wherein one or more virtual IP addresses are associated with the physical port, at least one of the one or more of the virtual IP addresses being an exclusive resource to a virtual server that includes as exclusive resources one or more file systems, wherein an exclusive resource of a virtual server can be seen by the virtual server and not by another virtual server;

determine whether any other physical port on the first node is good;

migrate all virtual IP addresses associated with the failed physical port to a good physical port on the first node if there is such a good port; and

migrate all virtual IP addresses associated with the failed physical port along with all virtual servers attached to such virtual IP addresses to a different, second node if there is no such good port on the first node.

Applicant : Gadir et al.
Serial No. : 09/911,902
Filed : July 23, 2001
Page : 10 of 15

Attorney's Docket No.: 07575-053001

35. (Currently Amended) A computer program product, tangibly stored on a computer-readable medium or propagated signal, for execution in a file server node in which one or more virtual servers each have one or more virtual IP addresses associated with physical ports, the product comprising instructions operable to cause a programmable processor to:

detect a failure of a physical port on a file server node, the node having two or more physical ports, the node having one or more virtual servers each ~~have~~having as exclusive resources one or more virtual IP addresses associated with physical ports and one or more file systems, wherein an exclusive resource of a virtual server can be seen by the virtual server and not by another virtual server;

identify one or more other physical ports on the file server node as being good; and
migrate each virtual IP addresses associated with the failed physical port to a good physical port on the file server node.

36. (Previously Presented) The product of claim 35, further comprising instructions to:

determine a load on each physical port on the first node; and
use the determined load for load balancing over the good physical ports when migrating the virtual IP addresses associated with the failed physical port to the good physical ports of the file server node.

37. (Previously Presented) The product of claim 35, wherein:

each physical port of the file server node is within a one of a plurality of subnets; and
virtual IP addresses are migrated preferentially to good physical port that is in the same subnet as the failed physical port.

Applicant : Gadir et al.
Serial No. : 09/911,902
Filed : July 23, 2001
Page : 11 of 15

Attorney's Docket No.: 07575-053001

38. (Currently Amended) A file server node, comprising:

two or more physical ports;

the node being configured to run two or more virtual servers, each virtual server having as exclusive resources a virtual interface to clients and one or more file systems, each virtual interface comprising a virtual IP address, wherein an exclusive resource of a virtual server can be seen by the virtual server and not by another virtual server;

the node being further configured to detect a failure of a first physical port, determine which other physical port or ports of the node is healthy, and to migrate all virtual IP addresses associated with the failed first physical port to a good physical port of the first node.

39. (Previously Presented) The file server node of claim 38, further configured to:

determine a load on each physical port; and

use the determined load for load balancing over the good physical ports when migrating the virtual IP addresses associated with the failed physical port to the good physical ports of the node.

40. (Previously Presented) The file server node of claim 38, wherein:

each physical port of the file server node is within a one of a plurality of subnets; and virtual IP addresses are migrated preferentially to good physical port that is in the same subnet as the failed physical port.